

Traction Adaptive Motion Planning and Control at the Limits of Handling

Downloaded from: https://research.chalmers.se, 2025-07-02 14:02 UTC

Citation for the original published paper (version of record):

Svensson, L., Bujarbaruah, M., Karsolia, A. et al (2022). Traction Adaptive Motion Planning and Control at the Limits of Handling. IEEE Transactions on Control Systems Technology, 30(5): 1888-1904. http://dx.doi.org/10.1109/TCST.2021.3129373

N.B. When citing this work, cite the original published paper.

© 2022 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, or reuse of any copyrighted component of this work in other works.

Traction Adaptive Motion Planning and Control at the Limits of Handling

Lars Svensson[®], *Member, IEEE*, Monimoy Bujarbaruah[®], Arpit Karsolia, Christian Berger[®], and Martin Törngren[®], *Senior Member, IEEE*

Abstract—In this article, we address the problem of motion planning and control at the limits of handling, under locally varying traction conditions. We propose a novel solution method where traction variations over the prediction horizon are represented by time-varying tire force constraints, derived from a predictive friction estimate. A constrained finite time optimal control problem (CFTOC) is solved in a receding horizon fashion, imposing these time-varying constraints. Furthermore, our method features an integrated sampling augmentation procedure that addresses the problems of infeasibility and sensitivity to local minima that arise at abrupt constraint alterations, for example, due to sudden friction changes. We validate the proposed algorithm on a Volvo FH16 heavy-duty vehicle, in a range of critical scenarios. Experimental results indicate that traction adaptive motion planning and control improves the vehicle's capacity to avoid accidents, both when adapting to low local traction, by ensuring dynamic feasibility of the planned motion, and when adapting to high local traction, by realizing high traction utilization.

Index Terms—Adaptive control, autonomous vehicles, collision avoidance, friction, motion planning, optimization-based motion planning, sampling-based motion planning, vehicle control.

I. INTRODUCTION

UTOMATED driving and advanced driver assistance technology show increasing potential to improve safety and mobility of transportation systems in the future. A major

Manuscript received 29 July 2021; accepted 28 October 2021. Date of publication 6 December 2021; date of current version 18 August 2022. This work was supported in part by the SAFER Open Research at AstaZero, in part by the Integrated Transport Research Lab at KTH, AutoDrive, under Agreement 737469, in part by the Intelligent Secure Trustable Things (InSecTT) under Agreement 876038, in part by the InSecTT (www.insectt.eu) through the ECSEL Joint Undertaking (JU), and in part by the JU through the European Union's Horizon 2020 Research and Innovation Programme and Austria, Sweden, Spain, Italy, France, Portugal, Ireland, Finland, Slovenia, Poland, The Netherlands, and Turkey. Recommended by Associate Editor A. Vahidi. (*Corresponding author: Lars Svenson.*)

Lars Svensson and Martin Törngren are with the Division of Mechatronics and Embedded Control Systems, KTH Royal Institute of Technology, 10044 Stockholm, Sweden (e-mail: larsvens@kth.se; martint@kth.se).

Monimoy Bujarbaruah is with the Model Predictive Control Lab, University of California at Berkeley, Berkeley, CA 94720 USA (e-mail: monimoyb@berkeley.edu).

Arpit Karsolia is with the Department of Electrical Engineering, Chalmers University of Technology, 41296 Gothenburg, Sweden (e-mail: arpit.karsolia@chalmers.se).

Christian Berger is with the Department of Computer Science and Engineering, University of Gothenburg, 40530 Gothenburg, Sweden (e-mail: christian.berger@gu.se).

Color versions of one or more figures in this article are available at https://doi.org/10.1109/TCST.2021.3129373.

Digital Object Identifier 10.1109/TCST.2021.3129373

challenge in driving automation is handling of critical situations, that is, situations that appear suddenly, in which an accident is imminent. Critical situations may originate from multiple sources, for example, unpredictable behavior of other traffic agents or rapid changes in the operational conditions. The steady progress in sensors as well as perception and prediction algorithms indicate that a subset of critical situations can be anticipated and avoided ahead of time. However, considering the inherent unpredictability of humans in the traffic environment and the wide range of operational conditions, it is unlikely that all critical situations can be anticipated and avoided ahead of time, without sacrificing efficiency and availability by enforcing overly cautious behavior. When a critical situation does occur, passenger comfort is no longer a priority, and if necessary, we wish to utilize the full physical capability of the vehicle to avoid the imminent accident. This motivates research in the field of motion planning and control at the limits of handling. Previous research efforts present methods for which performance approaches the physical limits [1], [2]. However, the motion capability of a road vehicle is greatly affected by locally varying traction conditions of the road. This phenomenon, which can be intuitively understood by the simple experiment depicted in Fig. 1, has been well studied and is included in standard vehicle dynamics literature [3], [4], but has only recently been considered in the context of motion planning at the limits of handling [5].

The work presented in this article represents a further step toward traction adaptive motion planning and control at the limits of handling. We propose a novel optimizationbased framework, in which the locally varying traction limit is represented as time-varying constraints on tire forces. The constraints are updated online as a function of the estimated friction coefficient and dynamic tire loads. Furthermore, we employ sampling augmentation to address the non-convexity and infeasibility issues of the arising optimization problem. We evaluate the proposed method in comparison with an equivalent non-adaptive scheme, in a sequence of experiments on various road surfaces.

The article presents an extended and refined version of the method that was originally introduced in [6]. The extension includes refinement of the mathematical framework, a more accurate model of the vehicle and its traction limitations, a novel, GPU accelerated method to generate feasible sample trajectories, and an extensive experimental evaluation.

This work is licensed under a Creative Commons Attribution 4.0 License. For more information, see https://creativecommons.org/licenses/by/4.0/



Fig. 1. Example of how operational conditions dictate the dynamic capabilities of a road vehicle. Entering the curve at 15 m/s, the vehicle comfortably negotiates the corner when the road is dry, but slides out into the opposing lane when the road is wet, causing a hazardous situation. (a) Dry road surface. (b) Wet road surface.

We summarize the contributions of the article as follows.

- 1) A real-time capable algorithm for traction adaptive motion planning and control that produces optimal solutions with respect to time-varying tire force constraints.
- 2) An experimental evaluation of traction adaptive motion planning and control, showing that the concept improves capacity to avoid accidents in a range of critical scenarios, when adapting to low as well as high local traction.

II. MOTIVATION

Critical scenarios at the limits of handling put the following high-level requirements on the local motion planning functionality.

- 1) A relatively complex vehicle model is required to accurately describe the vehicle in aggressive motions.
- A sufficiently long planning horizon is required to the handle the inertia and limited actuation capacity of the vehicle.
- 3) A short planning time is required in order to react promptly to a dynamically changing traffic scene.
- 4) The planner has to be able to make discrete decisions, for example, to go left or right of an obstacle.
- 5) Local traction variations should be considered.

Requirements 1–4 have been identified in previous academic works with state-of-the-art approaches presented in Section III. Next, we motivate why 5 should also be considered in this context.

From basic mechanics we have that for a single tire, the combined longitudinal, F_x , and lateral, F_y , force that can be exerted between tire and road is upper bounded by the friction coefficient μ times the normal force acting on the tire, F_z , that is,

$$\sqrt{F_x^2 + F_y^2} \le \mu F_z. \tag{1}$$



Fig. 2. Numerical example of how tire force limits of our test vehicle vary at different operational conditions and motion cases, here represented by the following parameters: Dry: $\mu = 0.8$, Wet: $\mu = 0.5$, Snow: $\mu = 0.3$, Accelerating: $a_x = 0.5 \text{ m/s}^2$, Constant speed: $a_x = 0 \text{ m/s}^2$, Braking: $a_x = -0.5 \text{ m/s}^2$.

This limit forms a circle in the lateral–longitudinal plane and is therefore commonly referred to as the friction circle.¹

Both μ and F_z typically vary substantially with time. The friction coefficient varies with, for example, road surface, tire temperature, and wear, and the normal force varies as a result of road inclination and dynamic forces caused by vehicle motion. Fig. 2 highlights the impact of this phenomenon in terms of motion planning near the physical limits. The figure visualizes to what extent the force limits of the front and rear tires vary under different operational conditions and motion cases.

The phenomenon presents an additional challenge to motion planning in critical situations. For example, not considering it and assuming static bounds on tire forces will have one of two consequences when planning aggressive evasive maneuvers.

- 1) *Planning of infeasible motions:* The real vehicle will be unable to track the plan.
- Reduced capability to avoid accident: The planner will not be able to fully utilize the available traction to avoid an accident.

Both 1 and 2 reduce the vehicle's capacity to avoid accidents in critical situations. This conclusion motivates us to pursue the development of algorithms for motion planning and control at the limits of handling, with a time-varying traction boundary. We refer to this concept as traction adaptive motion planning and control. Next, we give a brief summary of relevant academic works in related fields.

III. RELATED WORK

Motion planning algorithms for automated vehicles can be roughly classified into random sampling-based, graph

¹The friction circle is an approximation of the true force boundary of a tire. In reality, different tire–surface combinations will have differently shaped force boundaries. Therefore, the framework presented in Section V is designed to allow any convex shape to represent the force limit. In this study, however, we use a circular representation for simplicity.

search-based, trajectory roll-out, and optimization-based [7], [8]. Considering the special requirements in critical scenarios stated in Section II, there is an evident computational trade-off between properties 1, 2, and 3, and therefore, efficient solving of the motion planning problem is a key challenge. Previous studies, for example, [1], [2] show that optimization-based methods are well suited to handle such a computational tradeoff, but have less favorable properties for requirement 4, compared to the other three classes.

Random sampling-based and graph search-based algorithms are primarily designed to produce complex maneuver sequences [7], which is typically not required for short-term evasive maneuvers in critical situations. Also, such algorithms are inherently sequential, restricting efficient implementation through parallel computing. Therefore, we delimit our related work section to optimization-based and trajectory rollout algorithms, as well as a few relevant concepts from the fields of tire–road friction estimation and adaptive control.

A. Optimization-Based Methods

Optimization-based methods dominate the field of motion planning at the limits of handling. Compared to other methods, they are not restricted to a discretized search space, scale comparatively well with model complexity [9], and allow explicit representation of constraints on states and inputs.

The key concept of optimization-based motion planning is to solve a constrained finite time optimal control problem (CFTOC) at each planning iteration, where the solution is the control and state sequence that obeys the constraints and minimizes a cost function over a fixed time horizon. In general, the CFTOC representing a collision avoidance motion planning problem is nonconvex, due to the nonlinear vehicle dynamics and nonconvex state constraints representing obstacles [10]. The nonconvexity of the problem renders the problem nontrivial to solve under the time requirements associated with critical evasive maneuvers. However, there exist several methods such as quasi-Newton algorithms with line search or sequential quadratic programming (SQP) that are capable of solving the problem given a near-optimal initial guess of the optimal trajectory [11]-[13]. SQP is an iterative approach in which a quadratic, and hence convex, approximation of the problem is formed and solved at each iteration. The approximation is obtained by performing a second-order approximation of the cost function and a linearization of dynamics and constraints at the solution from the previous iteration. Thus, a quadratic program (QP) is solved at each iteration. Diehl et al. [14] proposed selecting the solution trajectory from the previous planning iteration as the initial guess and performed a single SQP iteration per planning iteration. This approach, called real time iteration (RTI), is exceptionally computationally efficient and has been proved successful in numerous applications, including motion planning for automated vehicles [1], [9], [15].

Ziegler *et al.* [9] highlighted a drawback of the RTI approach, noticing a sensitivity to local minima in situations where the motion planning problem contains discrete decision-making. The intuition behind the phenomenon is that at

each iteration, the optimizer searches for the best solution locally around the initial guess, and therefore it is unlikely to find a solution in a different part of the state space, which may contain the globally optimal solution. Liniger *et al.* [1] identify the same issue and mitigate it by means of a high-level path planner based on dynamic programming, a solution that works well in the particular application, but not for the general case.

B. Trajectory Roll-Out Methods

The trajectory roll-out concept represents a radically different approach to solving the motion planning problem. A large set of candidate trajectories are computed followed by a feasibility check and cost function evaluation that determines the optimal feasible candidate within the set [16]–[18]. Since the candidate trajectories are distributed across the drivable area, the method is less sensitive to local minima and thus performs well in situations where the motion planning problem contains discrete decisions.

Computing the individual trajectories connecting the initial state to a sampled goal state means solving a nonlinear optimal control problem. This is computationally intractable for large numbers of samples. However, for particular classes of trajectory candidates, for example, quintic and quartic polynomials, these problems have closed-form solutions [18]. While such a solution makes trajectory roll-out computationally tractable, the resulting reference trajectories are not suitable for maneuvers at the limits of handling, due to the inaccurate representation of vehicle dynamics and limitations.

C. Adaptive Control

To account for local variations in physical capabilities of the vehicle, we draw from developments in the field of adaptive control. Predictive control under model uncertainty is a well-studied topic [19]–[21]. Such frameworks allow the system to dynamically re-plan safer and more cost efficient trajectories with time, as model parameters get updated.

These and related concepts are increasingly being adopted to account for traction variations in motion planning and control of automated vehicles. For example, Alsterda *et al.* [5] present a novel approach in which a contingency plan, based on an alternative model parameter set, is computed in parallel to the nominal plan. As such, the method maintains a dynamically feasible contingency plan at all times, even in the event of, for example, a rapid friction reduction. In this work, we instead leverage the notion of adaptive control by updating model parameters based on a predictive friction estimate.

D. Predictive Tire–Road Friction Estimation

There exists a large body of research on the topic of tire–road friction estimation from on-board vehicle measurements [22], [23]. Such estimates have previously been shown to improve performance in several vehicle control applications, for example, automatic emergency braking [24], [25] and path tracking control at deteriorated traction [26], [27]. With the recent developments in computer vision and machine learning,

methods for predictive friction estimation using camera images have emerged [28]–[32]. A predictive friction estimate enables motion planning according to the traction conditions ahead of the vehicle. This could also be obtained using a classical method such as [22], [23], under the assumption that the traction conditions are unchanged for a short distance ahead of the vehicle.

To perform motion planning and control with respect to such predictive friction estimates, the planner must reliably handle situations where the friction estimate varies over the prediction horizon or changes between subsequent planning iterations.

E. Summary

To summarize, optimization-based methods (RTI in particular) provide good performance in handling the trade-off between requirements 1, 2, and 3 in Section II, but struggle with 4. Trajectory roll-out methods, on the other hand, inherently handle 4, but have much lower performance in the trade-off between 1, 2, and 3. The method proposed in this work combines RTI, trajectory roll-out, and adaptive control to collectively address requirements 1–5. We leverage the notion of adaptive control by utilizing updated model information derived from a predictive friction estimate,² to account for time-varying traction limitations in local motion planning, including discrete decisions.

IV. PROBLEM FORMULATION

Our goal in algorithm design is to develop a method that efficiently computes dynamically feasible reference trajectories at a high degree of traction utilization on varying road surfaces, based on a predictive friction estimate. Furthermore, the algorithm must handle cases where the predictive friction estimate changes over the course of the planning horizon and/or changes abruptly between planning iterations. In this section, we introduce the mathematical details of the considered motion planning problem and state the challenges it presents in further detail.

A. Planning Model

We derive a continuous time vehicle model with timevarying parameters, as shown in Appendix B, and denote it $f_c(\cdot, \cdot)$. We let t_c denote continuous time. For controller design, we obtain a discrete time version of this model as

$$f(x_t, u_t) = x_t + T_s f_c(x(tT_s), u(tT_s))$$
(2)

with time index t, time step T_s , and the state and input vectors

$$x = [s, d, \Delta \psi, \dot{\psi}, v_x, v_y]^\top \in \mathbb{R}'$$
$$u = [F_{yf}, F_{xf}, F_{xr}]^\top \in \mathbb{R}^m.$$

Remark 1: We avoid separate notations for "nominal" states from (2) and the true states from the vehicle for the sake of notational simplicity in the remainder of the article. This can

be done as we do not opt for a robust control design approach, such as [33]–[35].

State variables *s*, *d*, and $\Delta \psi$ represent the progress along, the lateral deviation from, and relative orientation to the lane center, respectively. Variables $\dot{\psi}$, v_x , and v_y denote yaw angular velocity, longitudinal, and lateral velocity of the vehicle, respectively. We employ the single-track modeling technique [4], in which the combined influence of two wheels on an axle is represented by a single tire in the middle of the axle and select as control inputs: the lateral force on the front tire, F_{yf} , the longitudinal force on the front tire, F_{xf} , and the longitudinal force on the rear tire F_{xr} .

Remark 2: The primary reason for selecting tire force control inputs is that it allows us to explicitly model the locally varying traction limit as a time-varying constraint on the inputs. Furthermore, it contributes to the generality of the solution, since it can be instantiated on a wide range of vehicle types by setting a small number of easily measured parameters. A drawback compared to selecting steering angle as the control input for the lateral dimension is that it is non-trivial to represent actuator limitations such as maximum steering angle. However, for the speed range associated with the type of maneuvers, we are targeting ($v_x > 5$ m/s) this problem does not occur. In order to translate the resulting tire force commands into the vehicle's actual control inputs, a vehicle specific control interface is required. The control interface developed for our test vehicle is described in Appendix C.

In addition to the dynamics in the *xy*-plane, the model includes relations to determine the front and rear normal forces F_{zf} and F_{zr} from the state x_t . This, together with a predictive friction estimate, allows us to obtain a time-varying traction force boundary through (1). Further details on the model are provided in Appendix B.

B. Optimal Control Problem

In optimization-based motion planning, a CFTOC is solved at regular time intervals, producing a new motion plan at each planning iteration as its solution. For the motion planning problem with traction limit variation, we consider the CFTOC

$$\begin{array}{ll} \min_{u_{0|t},\dots,u_{N-1|t}} & J(\mathcal{T}_{t}) \\ \text{s.t.,} & x_{k+1|t} = f\left(x_{k|t}, u_{k|t}\right) \\ & u_{k|t} \in \mathcal{U}_{k|t} \\ & x_{k|t} \in \mathcal{X}_{k|t} \\ & \forall k \in \{0, 1, \dots, (N-1)\} \\ & x_{0|t} = x_{t}, \ x_{N|t} \in \mathcal{X}_{k|t} \end{array} \tag{3}$$

where the variables $\{u_{0|t}, u_{1|t}, \ldots, u_{N-1|t}\}\$ denote a predicted input sequence at time step *t*, which, when applied through the vehicle model $f(\cdot, \cdot)$, yields the corresponding predicted state sequence $\{x_{0|t}, x_{1|t}, \ldots, x_{N|t}\}\$. The closed-loop state is reset again when x_{t+1} is measured. We denote a predicted trajectory consisting of the associated predicted state and control sequences as

$$\mathcal{I}_t = \left\{ \{ x_{k|t} \}_{k=0}^N, \{ u_{k|t} \}_{k=0}^{N-1} \right\}$$

 $^{^{2}}$ We designed the method to receive a predictive friction estimate from a state-of-the-art method, but the estimation itself is not included in the scope of this article.

which is used in the selected quadratic cost function

$$J(\mathcal{T}_{t}) = x_{N|t}^{\top} Q_{N} x_{N|t} + \sum_{k=0}^{N-1} (x_{k|t}^{\top} Q x_{k|t} + u_{k|t}^{\top} R u_{k|t})$$
(4)

where $Q_N > 0$, Q > 0, and R > 0 are weight matrices that can be tuned by the designer. The choice here of a quadratic cost in (4) is done without loss of generality. Standard methods exists to make local quadratic approximations of non-quadratic cost functions (see Section III-A).

In (3), the state trajectory constraints $\mathcal{X}_{k|t}$, for all $k \in \{0, 1, ..., N\}$, represent collision-free configurations of the vehicle with respect to the road geometry, static obstacles, as well as positions and predicted positions of dynamic obstacles.

Notice that the input constraints $\mathcal{U}_{k|t}$ are time-varying along the prediction horizon for $k \in \{0, 1, ..., (N-1)\}$. This is the mechanism we use to adapt the planned motion to local traction limitations. The set $\mathcal{U}_{k|t}$ is computed at run time as a function of the predicted friction estimate $\mu_{k|t}$ and normal forces F_{zf} and F_{zr} . A further description on how the tire force limits are obtained is described in Appendix B. Finally, we let $\mathcal{T}_t^* = \{\{x_{k|t}^*\}_{k=0}^N, \{u_{k|t}^*\}_{k=0}^{N-1}\}$ denote the optimal predicted trajectory.

C. Applying a State-of-the-Art Approach: RTI

The RTI approach [14], [36] is a state-of-the-art approach for solving CFTOCs that enables extraordinary computational efficiency [37]. In the context of motion planning in critical situations, computational efficiency is essential to achieving an acceptable trade-off between properties 1, 2, and 3 of Section II. Hence, we select RTI as a starting point for algorithm development.

The central idea of the algorithm is to use the solution from the previous planning iteration to formulate a QP approximation of the CFTOC, which can be solved efficiently. The quadratic formulation is obtained by linearizing the dynamics and expressing the constraints as linear inequalities [10]. At a single planning iteration starting at time t, a standard RTI algorithm goes through the following steps.

First, the solution from the previous iteration $\mathcal{T}_{t-1}^{\star}$ is shifted one step forward in time, to compensate for the time passing between planning iterations. Thus, at time step *t*, an initial guess is obtained as

$$\hat{\mathcal{T}}_t = \left\{ \{ \hat{x}_{k|t} \}_{k=0}^N, \{ \hat{u}_{k|t} \}_{k=0}^{N-1} \right\}$$

where $\hat{x}_{k|t} = x_{k+1|t-1}^{\star}$ and $\hat{u}_{k|t} = u_{k+1|t-1}^{\star}$ for $k \in \{1, 2, \dots, (N-1)\}$.³ Next, the dynamics are linearized about \hat{T}_t . The system model matrices are obtained as

$$A_{k|t} = \frac{\partial f}{\partial x} \bigg|_{(\hat{x}_{k|t}, \hat{u}_{k|t})}, \quad B_{k|t} = \frac{\partial f}{\partial u} \bigg|_{(\hat{x}_{k|t}, \hat{u}_{k|t})}$$

³Forward shifting is not possible for the *N*th predicted state. So, the final control input is duplicated $\hat{u}_{N-1|t} = u_{N-1|t-1}^*$ and the final state is obtained by integrating the dynamics forward $\hat{x}_{N|t} = f(x_{N-1|t-1}^*, u_{N-1|t-1}^*)$.

for all $k \in \{0, 1, ..., (N-1)\}$. State and input constraints are expressed as sets of linear inequalities

$$\mathcal{P}_{k|t}^{\mathcal{X}} = \{x : H_{k|t}^{\mathcal{X}} x \le h_{k|t}^{\mathcal{X}}\} \subseteq \mathcal{X}_{k|t} \\ \mathcal{P}_{k|t}^{\mathcal{U}} = \{u : H_{k|t}^{\mathcal{U}} u \le h_{k|t}^{\mathcal{U}}\} \subseteq \mathcal{U}_{k|t}$$

with $H_{k|t}^{\mathcal{X}} \in \mathbb{R}^{s_x \times n}$, $h_{k|t}^{\mathcal{X}} \in \mathbb{R}^{s_x}$, for all $k \in \{0, 1, \dots, N\}$, and $H_{k|t}^{\mathcal{U}} \in \mathbb{R}^{s_u \times m}$, $h_{k|t}^{\mathcal{U}} \in \mathbb{R}^{s_u}$ for all $k \in \{0, 1, \dots, (N-1)\}$. The procedure we use to linearize the input constraints of our vehicle model is described in Appendix B.

The final RTI style reformulation of (3) around an initial guess \hat{T}_t is shown in the following equation:

$$\begin{array}{l} \min_{\Delta u_{0|t},...,\Delta u_{N-1|t}} \quad J(\mathcal{T}_{t}), \text{ with } \mathcal{T}_{t} = \{\{x_{k|t}\}_{k=0}^{N}, \{u_{k|t}\}_{k=0}^{N-1}\} \\
\text{s.t.}, \quad x_{k+1|t} = A_{k|t}(\Delta x_{k|t}) + B_{k|t}(\Delta u_{k|t}) + \hat{x}_{k+1|t}^{*} \\
\quad (\text{linearized about } \hat{\mathcal{T}}_{t}) \\
u_{k|t} \in \mathcal{P}_{k|t}^{\mathcal{U}} \\
\forall k \in \{0, 1, \dots, (N-1)\}, \text{ with} \\
\left[\Delta x_{k|t}, \Delta u_{k|t}\right] = [x_{k|t} - \hat{x}_{k|t}', u_{k|t} - \hat{u}_{k|t}'] \\
x_{k|t} \in \mathcal{P}_{k|t}^{\mathcal{X}} \\
\forall k \in \{0, 1, \dots, N\} \\
x_{0|t} = x_{t}.
\end{array}$$
(5)

When a solution of (5) for time step t is obtained, the control

$$u_t = u_{0|t}^{\star}$$

is applied to the vehicle. At the next time step, (t + 1), Problem (5) is solved again.

D. Challenges in Solving the RTI Problem (5)

The RTI algorithm builds on the assumption that $\mathcal{T}_{t-1}^{\star}$ is a feasible and near-optimal initial guess for time step *t*. If this is not the case, performance of the algorithm deteriorates, leading to three essential drawbacks with respect to the problem formulation of this article.

- Infeasibility w.r.t. P^U_{k|t}: In the presence of time-varying input constraints, there is no guarantee that at time step t the input sequence from, T^{*}_{t-1} is feasible with respect to P^U_{k|t}, that is, the condition u^{*}_{k|t-1} ∈ P^U_{k|t} may or may not hold for all k ∈ {1, 2, ..., N}.
 Infeasibility w.r.t. P^X_{k|t}: Irrespective of 1, in a dynami-
- 2) Infeasibility w.r.t. $\mathcal{P}_{k|t}^{\mathcal{X}}$: Irrespective of 1, in a dynamically changing environment with sensor limitations and unpredictably moving human actors, there is no guarantee that at time step t, the state sequence of $\mathcal{T}_{t-1}^{\star}$ will be collision-free, that is, the condition $x_{k|t-1}^{\star} \in \mathcal{P}_{k|t}^{\mathcal{X}}$ may or may not hold for all $k \in \{1, 2, ..., N\}$.
- 3) Sensitivity to local minima: Due to the fact that the search for \mathcal{T}_t^* is only performed locally around \mathcal{T}_{t-1}^* , RTI is prone to suboptimal decision-making where the planning problem includes discrete decision making. In realistic traffic, new discrete decisions may appear dynamically [9].

Problem 1 is unique to our problem formulation and stems from the addition of traction adaptation through time-varying input constraints $\mathcal{P}_{klt}^{\mathcal{U}}$ introduced in (3). A naive solution is to project $u_{k|t-1}^{\star}$ onto $\mathcal{P}_{k|t}^{\mathcal{U}}$ for all $k \in \{1, 2, \dots, (N-1)\}$ and recompute \hat{T}_t from the dynamics model (this is the first step of our proposed algorithm, see Section V-A). However, this is not a complete solution as it, for non-trivial cases, merely transforms Problem 1 into Problem 2. Problem 2 has been identified in previous works and is unavoidable in a dynamic traffic environment. Computational feasibility can be obtained by introducing slack variables, such that state constraint violations are allowed, but with a very high cost [1], [38]. Commonly, however, this approach merely transforms Problem 2 into Problem 3. Problem 3 in turn may be handled for specific applications by a higher-level decisionmaking layer [1]. In critical situations, however, the optimal decision, for example, to swerve left or right of a suddenly appearing obstacle, is highly dependent on, for example, the geometry of the traffic scene, the vehicle's state, dynamics, and dynamic limitations. This information is typically only available at the local motion planning level and, therefore, an informed decision can only be made there. In such critical decision-making, insufficient scene understanding may lead to suboptimal decisions regarding the selected maneuvers, which in turn reduces capacity to avoid accidents. Trajectory roll-out methods, in contrast to optimization-based methods, inherently handle such decisions by construction at the local planning level, as described in Section III-B.

In Section V, we present a novel variant of RTI, which we refer to as sampling augmented adaptive RTI (SAA-RTI), to jointly address 1, 2, and 3.

V. SAMPLING AUGMENTED ADAPTIVE RTI

The SAA-RTI algorithm decomposes the solving of (3) into four steps using the same planning horizon N and time discretization step T_s . Furthermore, steps C and D use the same cost function (4). The steps are:

- 1) ensuring input feasibility of $\mathcal{T}_{t-1}^{\star}$;
- generating additional initial guess candidates through sampling;
- 3) selection of initial guess trajectory;
- 4) constraint adaptive trajectory optimization.

Together, steps A–D remedy the limitations 1, 2, and 3 of RTI introduced in Section IV-D. In Sections V-A–V-D, we explain each step in further detail.

A. Ensuring Input Feasibility of T_{t-1}^{\star}

Due to the time-varying input constraints, situations when $u_{k|t-1}^* \notin \mathcal{P}_{k|t}^{\mathcal{U}}$ for $k \in \{1, 2, ..., (N-1)\}$ may occur. To resolve this, we replace the shifting procedure in conventional RTI as

$$\tilde{u}_{k|t} = \operatorname{Proj}_{\mathcal{P}_{kk}^{\mathcal{U}}}(u_{k+1|t-1}^{\star}) \quad \forall k \in \{0, 1, \dots, (N-1)\}$$
(6)

where $\operatorname{Proj}_{\Omega}(\cdot)$ denotes Euclidean projection operation to the set Ω . Thus, we modify individual inputs if needed, such that $\tilde{u}_{k|t} \in \mathcal{P}_{k|t}^{\mathcal{U}}$ for $k \in \{0, 1, \ldots, (N-1)\}$. Then, under the new input sequence, we propagate the system forward from $\tilde{x}_{0|t} = x_{1|t-1}^*$ to determine the state sequence $\tilde{x}_{k+1|t} = f(\tilde{x}_{k|t}, \tilde{u}_{k|t}), k \in \{1, 2, \ldots, (N-1)\}$. Therefore, we obtain an initial guess

$$\tilde{\mathcal{T}}_t = \left\{ \{ \tilde{x}_{k|t} \}_{k=0}^N, \{ \tilde{u}_{k|t} \}_{k=0}^{N-1} \right\}$$
(7)

that satisfies dynamics and time-varying input constraints, resolving limitation 1 of Section IV.

We acknowledge that the procedure may alter the state sequence such that state constraints are violated, that is, $\tilde{x}_{k|t} \notin \mathcal{P}_{k|t}^{\mathcal{X}}$, for some $k \in \{0, 1, ..., N\}$. However, as stated in Problem 2 in Section IV-D, this problem has to be addressed regardless of this addition, due to the dynamic environment. We address the problem in the next step of the algorithm by employing the trajectory roll-out concept [16].

B. Generating Additional Initial Guess Candidates Through Sampling

To handle the case when the initial guess (7) violates the state constraints, $\tilde{x}_{k|t} \notin \mathcal{P}_{k|t}^{\mathcal{X}}$ for any $k \in \{1, 2, ..., N\}$, that is, Problem 2, and to avoid highly suboptimal local minima, that is, Problem 3, we employ a further developed version of the sampling augmentation concept that was first introduced in [6].

The procedure starts by selecting N_s reference states $x_{ref}^{(i)}$, $i \in \{1, 2, ..., N_s\}$. Reference values for d and v_x are uniformly distributed within admissible values, the *s*-reference selected as the current *s* of the vehicle, and remaining state references are set to zero. For each reference state, we integrate the dynamics forward

$$\hat{x}_{k+1|t}^{(i)} = f\left(\hat{x}_{k|t}^{(i)}, \hat{u}_{k|t}^{(i)}\right) \quad \forall k \in \{0, 1, \dots, (N-1)\}$$
(8)

from the initial state $\hat{x}_{0|t}^{(i)} = x_t$, under a linear quadratic tracking (LQT) controller, for all $i \in \{1, 2, ..., N_s\}$, that minimizes deviation from the reference state. At each time step k, the optimal control input is computed as

$$\hat{u}_{k|t}^{(i)} = -K_t \left(\hat{x}_{k|t}^{(i)} - x_{\text{ref}}^{(i)} \right) \tag{9}$$

where K_t is the optimal gain matrix, recomputed at each planning iteration by linearizing the dynamics around the initial state x_t and solving the discrete-time algebraic Riccati equation. The cost function weights are selected such that behavior is primarily dominated by errors in d and v_x . Errors, for example, in s, have small influence.

To ensure $\hat{u}_{k|t}^{(i)} \in \mathcal{P}_{k|t}^{\mathcal{U}}$ for $k \in \{1, 2, ..., N\}$ we apply the procedure described in Section V-A to project all control inputs onto the feasible set before propagating the state. The N_s candidate trajectories

$$\hat{\mathcal{T}}_{t}^{(i)} = \left\{ \{ \hat{x}_{k|t}^{(i)} \}_{k=0}^{N}, \{ \hat{u}_{k|t}^{(i)} \}_{k=0}^{N-1} \right\}$$

for all $i \in \{1, 2, ..., N_s\}$ are stored in the trajectory set

$$\hat{\mathcal{S}}_t = \bigcup_{i=1}^{N_s} \hat{\mathcal{I}}_t^{(i)}.$$
(10)

Notice that following the computation of K_t , each trajectory is computed independently, enabling efficient GPU acceleration.

C. Selection of Initial Guess Trajectory

To determine the best available initial guess, each candidate trajectory in the set $\hat{S}_t \cup \tilde{T}_t$ is evaluated based on the shared

cost function $J(\cdot)$ in (4). The lowest cost candidate, that is,

$$\hat{\mathcal{T}}'_t = \arg\min_{\hat{\mathcal{T}}_t \in \hat{\mathcal{S}}_t \cup \tilde{\mathcal{T}}_t} (J(\hat{\mathcal{T}}_t))$$
(11)

is selected for subsequent optimization. Trajectory candidates that violate state constraints, that is, $\hat{x}_{k|t}$ ∉ $\mathcal{P}_{k|t}^{\mathcal{X}}$ for any $k \in \{1, 2, \dots, N\}$ are disregarded in the selection.

In case $\tilde{\mathcal{T}}_t$ is colliding or is in a highly suboptimal local minima, an alternative feasible and near-optimal initial guess can be obtained from \hat{S}_t . Thus, we exploit the inherent redundancy in the trajectory roll-out concept [18], to resolve Problems 2 and 3. This property of the algorithm is experimentally verified in Section VI-D.

Remark 3: In cases when changes to the time-varving state and input constraints are small between planning iterations, it is highly probable that $\hat{T}'_t = \tilde{T}_t$. In this case, algorithm behavior is identical to that of RTI. This is further elaborated in Section VI-D.

D. Constraint Adaptive Trajectory Optimization

With the initial guess \hat{T}'_t obtained from (11), the final step of the SAA-RTI algorithm is to solve the OP

$$\begin{array}{ll} \min_{\Delta u_{0|t},...,\Delta u_{N-1|t}} & J(\mathcal{T}_{t}), \text{ with } \mathcal{T}_{t} = \{\{x_{k|t}\}_{k=0}^{N}, \{u_{k|t}\}_{k=0}^{N-1}\} \\ \text{s.t.}, & x_{k+1|t} = A_{k|t}(\Delta x_{k|t}) + B_{k|t}(\Delta u_{k|t}) + \hat{x}'_{k+1|t} \\ & (\text{linearized about } \hat{\mathcal{T}}'_{t}) \\ & u_{k|t} \in \mathcal{P}_{k|t}^{\mathcal{U}} \\ & \forall k \in \{0, 1, \dots, (N-1)\}, \text{ with} \\ & [\Delta x_{k|t}, \Delta u_{k|t}] = [x_{k|t} - \hat{x}'_{k|t}, u_{k|t} - \hat{u}'_{k|t}] \\ & x_{k|t} \in \mathcal{P}_{k|t}^{\mathcal{X}} \\ & \forall k \in \{0, 1, \dots, N\} \\ & x_{0|t} = x_{t}. \end{array} \tag{12}$$

In contrast to (5), the QP is initialized at \hat{T}'_t , that is, with the system model matrices given as

$$A_{k|t} = \frac{\partial f}{\partial x} \bigg|_{(\hat{x}'_{k|t}, \hat{u}'_{k|t})}, \quad B_{k|t} = \frac{\partial f}{\partial u} \bigg|_{(\hat{x}'_{k|t}, \hat{u}'_{k|t})}$$

for all $k \in \{0, 1, \dots, (N-1)\}$. After solving (12) at any time step t, we apply the closed-loop command

$$u_t = u_{0|t}^{\star} \tag{13}$$

to the vehicle system through the low-level control interface described in Appendix C. At time step (t + 1), (12) is re-solved.

E. Summary and Implications for Traction Adaptation

We summarize SAA-RTI in Algorithm 1. Here, M_t represents local map features, for example, lane boundaries and static obstacles, and O_t denotes dynamic obstacles at time step t. Note that we assume availability of a predictive tire-road friction estimates $\mu_{k|t}$ for $k \in \{0, 1, \dots, (N-1)\}$ at any time step t.

Algorithm 1 The SAA-RTI Algorithm

1: Initialization: $\mathcal{T}_{t-1}^{\star} = \emptyset$ 2: while t > 0 do

Inputs: $x_t, \mathcal{T}_{t-1}^{\star}, M_t, O_t, \mu_{k|t}, \forall k \in \{0, 1, \dots, N\}$ 3:

4: **if**
$$\mathcal{T}_{t-1}^{\star} \neq \emptyset$$
 then

5: $\tilde{\mathcal{T}}_t \leftarrow \text{ShiftAndEnsureFeasibility}(\mathcal{T}_{t-1}^{\star}, \mu_{k|t}, \forall k)$

- $\tilde{\mathcal{T}}_t = \emptyset$ 7:
- end if 8:
- $\hat{\mathcal{S}}_t \leftarrow \text{FeasibleTrajRollout}(x_t, M_t, \mu_{k|t}, \forall k) \text{ from (10)}$ 9:
- 10: $\hat{\mathcal{T}}'_t = \arg\min_{\tilde{\mathcal{T}}} (J(\hat{\mathcal{T}}_t)) \text{ from (11)}$

11:
$$A_{k|t}, B_{k|t} \leftarrow \text{LinearizeModel}(\hat{\mathcal{T}}'_t, \forall k)$$

12:
$$\mathcal{P}_{k|t}^{\mathcal{U}} \leftarrow \text{InputConstraints}(\hat{\mathcal{T}}_{t}', \mu_{k|t}, \forall k)$$

 $\mathcal{P}_{k|t}^{\mathcal{X}} \leftarrow \text{InputConstraints}(T_t', \mu_k|t, \forall k)$ $\mathcal{P}_{k|t}^{\mathcal{X}} \leftarrow \text{StateConstraints}(\hat{T}_t', O_t, M_t, \forall k)$ 13:

14:
$$\mathcal{T}_{t}^{\star} = \{\{x_{h|t}^{\star}\}_{k=0}^{N}, \{u_{h|t}^{\star}\}_{k=0}^{N-1}\} \leftarrow \text{Solve (12)}$$

5: Apply
$$u_{t|t}^{\star}$$
 to the vehicle

16: end while

1



Fig. 3. Example of planned behavior when the friction coefficient varies over the prediction horizon. The proposed approach plans dynamically feasible motions at the (varying) limit of tire adhesion. (a) Position and orientation of planned motion. (b) Velocity and front tire force of planned motion.

The proposed approach of modeling locally varying tire force limits as time-varying input constraints allows the vehicle to plan motions that optimally utilizes traction throughout the prediction horizon. Fig. 3 shows an example where μ drops from 0.8 to 0.3 in the middle of the horizon. Recall that t_c denotes our continuous time variable. The planned motion exploits the initial high traction to reduce velocity [Fig. 3(b)], and position the vehicle such that the curvature of the trajectory is reduced over the low traction part of the corner [Fig. 3(b)]. This is done in a coordinated manner such that planned tire forces, shown for the front tire in Fig. 3(b), are at the (varying) limit throughout the planned maneuver, that is, for the front tire, $\sqrt{(F_{xf})_{k|t}^2 + (F_{yf})_{k|t}^2} = \mu_{k|t}(F_{zf})_{k|t}$.

The SAA-RTI algorithm ensures efficient and reliable solving of the traction adaptive motion planning problem. The sampling augmentation procedure ensures that the optimization procedure is always provided a feasible and near-optimal



Fig. 4. Initial guess candidates when estimated traction changes abruptly between planning iterations. Sampling augmentation mitigates the problems of using the solution from the previous iteration as initial guess. (a) Position and orientation of initial guess candidates and optimized solution. (b) Velocity state of initial guess candidates and optimized solution.

initial guess, even when the predictive friction estimate changes dramatically between planning iterations. Fig. 4 shows an example where the estimate of μ has dropped from 0.8 to 0.3 between the previous and the current planning iteration. The figure shows position and orientation, Fig. 4(a), and velocity states, Fig. 4(b), of the forward shifted and input feasibility adjusted solution from the previous iteration $\tilde{\mathcal{T}}_t$ (blue), the lowest cost SAA-RTI initial guess candidate \hat{T}_t (orange) and the final optimized trajectory \mathcal{T}_t^{\star} (green). It is evident from Fig. 4(a) that even when adjusted for input feasibility, Problem 1, $\tilde{\mathcal{T}}_t$ is still not feasible with respect to state constraints, Problem 2. The addition of sampling augmentation reliably alleviates this problem by providing an abundance of additional feasible candidates and selecting the best available option $\hat{\mathcal{T}}'_t$. In addition, since the candidates are spread across the drivable area, sampling augmentation greatly reduces sensitivity to local minima in the presence of discrete decisions, Problem 3, which we will revisit in Section VI-D.

VI. EXPERIMENTAL EVALUATION

We evaluate the performance of traction adaptive motion planning and control, realized by Algorithm 1, in terms of its capacity to avoid accidents in a set of critical scenarios. We compare the traction adaptive algorithm (abbreviated TA) with a baseline non-adaptive RTI scheme with static tire force constraints (abbreviated NA). The TA scheme uses time-varying tire force constraints computed as a function of predicted normal forces and a predictive tire–road friction estimate $\mu_{k|t} = \mu_{est}(s)$ for all $k \in \{0, 1, \dots, (N-1)\}$ as per (16) and (17) in Appendix B with traction utilization factor $\lambda = 0.90$ for all scenarios. The NA scheme, on the other hand, uses static tire force constraints associated with a static friction estimate $\mu_{k|t} = \mu_{sta}$ for all $k \in \{0, 1, \dots, (N-1)\}$ and $a_x = 0$. All evaluated configurations use N = 40 and $T_s = 0.1$ s, that is, a planning horizon length of 4 s. The critical scenarios selected for evaluation are:

- 1) Turn at low local traction;
- 2) Collision avoidance at low local traction;
- 3) Collision avoidance at high local traction;
- 4) Collision avoidance with discrete decisions.

Scenarios 1–3 are selected to test the accident avoidance performance impact of traction adaptation, whereas Scenario 4 highlights the impact of avoiding local minima in the planning problem.

For the evaluation in this article, we assume that an accurate predictive friction estimate is available, that is, $\mu_{est}(s) = \mu_{gt}(s)$, where $\mu_{gt}(s)$ denotes the ground-truth tire–road friction, which was obtained for the experiments by performing friction measurements ahead of time and storing the values in a map of the track. The tests were performed on flat and level road surfaces and therefore, evaluation of the effects of road inclination is not included in this work, that is, $\theta = \phi = 0$ rad. in (14), Appendix B. In the collision avoidance scenarios, all suddenly appearing obstacles remain at a static position after appearing.

Remark 4: This delimitation is not done because of limitations in the evaluated algorithms with respect to dynamic obstacles. For both the TA and NA schemes, state constraints $\mathcal{P}_{k|t}^{\mathcal{X}}$ are set individually for each $k \in \{0, 1, \ldots, N\}$. Hence, including avoidance of, for example, a confidence interval of predicted positions of the obstacle, is trivial from the planning perspective. Rather, the delimitation was made to ensure consistency between runs and to avoid introducing additional uncertainty from the prediction functionality. Since the traction adaptive and the non-adaptive schemes are equivalent in this regard, we are confident that the conclusions we make for static obstacles will also hold for dynamic obstacles.

In Sections VI-A–VI-D, we present and discuss results from each of the scenarios. The experimental setup is detailed in Appendix A.

A. Scenario 1: Turn at Low Local Traction

At the start of the first critical scenario at $t_c = 0$ s and s = 0 m, the vehicle approaches a turn at 8 m/s. Traction has deteriorated such that $\mu_{gt}(s) = 0.2$ for $s \ge 0$ m instead of the previous $\mu_{gt}(s) = 0.8$ for s < 0 m. The objective of the vehicle is to maintain its initial velocity and stay close to the center of the lane. Fig. 5 shows a driving behavior analysis comparing the TA scheme with time-varying tire force constraints corresponding to $\mu_{est}(s) = \mu_{gt}(s)$ and the NA scheme with static tire force constraints corresponding to $\mu_{sta} = 0.8$.

1) Not Adapting to Low Local Traction: We make the following observations of the behavior under the NA scheme. At $t_c = 0$ s, before reaching the corner, Fig. 5(a), the planned motion appears to follow the intended objective of tracking the center of the lane and maintaining velocity. Immediately after entering the corner, Fig. 5(c), the vehicle starts to exhibit under-steering behavior that was not previously anticipated, that is, the planned motion at $t_c = 0$ s did not include this motion, and the plan gets gradually shifted toward the outside of the corner. The vehicle continues to under-steer heavily [see



Fig. 5. Scenario 1, turn at low μ . (a)–(h) Overhead visualizations of the planned motion. The left and right columns correspond to the non-adaptive (NA) and traction adaptive (TA) schemes, respectively. (i) Planned tire forces and prevailing force limits for the front tires at $t_c = 0$ s. (j)–(m) Comparisons in terms of deviation from the lane center, velocity, front wheel slip angle, and lateral acceleration. (a) NA, $t_c = 0$ s. (b) TA, $t_c = 0$ s. (c) NA, $t_c = 2.5$ s (under-steering). (d) TA, $t_c = 2.5$ s. (e) NA, $t_c = 3.5$ s (exiting lane). (f) TA, $t_c = 4$ s. (g) NA, $t_c = 7$ s (accident). (h) TA, $t_c = 5$ s (no accident). (i) Planned tire forces at $t_c = 0$ s [corresponding to (a) and (b)]. (j) Position w.r.t. lane center. (k) Forward velocity. (l) Front wheel slip angle. (m) Lateral acceleration.

front slip angle in Fig. 5(1)], such that it eventually enters the opposing lane, Fig. 5(e), at a velocity of 7 m/s Fig. 5(k). At this point in the experiment, the planner disengages and full braking is applied. The vehicle eventually comes to a stop 15 m outside of the lane, Fig. 5(g). The planned trajectory at the beginning of the scenario, Fig. 5(a), differs dramatically from the resulting closed-loop trajectory, Fig. 5(c), (e), and (g).

The reason behind this undesirable behavior is evident from the planned tire forces at $t_c = 0$ s, plotted in Fig. 5(i). The dashed blue and dotted orange lines represent the maximum horizontal forces that can be generated by the front and rear tires, respectively, according to (1), given the recorded state trajectory. It is clear that the non-adaptive scheme plans tire forces (solid blue) that violate the physical limitations. Therefore, the vehicle is unable to realize the planned motion and promptly veers out of its lane.

2) Adapting to Low Local Traction: The adaptive scheme, however, sets the tire force constraints dynamically according to the predictive friction estimate $\mu_{est}(s)$. Therefore, the planned tire forces [dashed-dotted orange in Fig. 5(i)] do not violate the physical limitations of the tires. Instead, it plans a motion, Fig. 5(b), that minimizes the deviation from the objective (track the center of the lane and maintain 8 m/s), while satisfying the updated tire force constraints. With the TA scheme, the vehicle slows down from 8 to 5.5 m/s before entering the corner, Fig. 5(k) and keeps a position on the outside of the turn. About a third of the way through the turn, Fig. 5(d), the vehicle has cut across the lane and is positioned along the inside of the corner, Fig. 5(j). Toward the end of the turn, Fig. 5(h), it starts cutting back across to the outside of the corner. Finally, the vehicle completes the turn safely, positioned at the outer edge of the lane, Fig. 5(h) and (j). In contrast to the NA scheme, here the planned trajectory at the beginning of the scenario, Fig. 5(b), corresponds well with the closed-loop trajectory of the vehicle, Fig. 5(d), (f), and (h).

We break this behavior down into three coordinated characteristics of the maneuver that enable the vehicle to successfully negotiate the turn. First, contrary to the NA scheme, the TA scheme generates feasible tire force commands, utilizing approximately 90% ($\lambda = 0.90$) of the available traction for this surface [see Fig. 5(i)]. This in turn generates slip angles within acceptable limits, Fig 5(1), yielding high lateral acceleration, Fig 5(m). Second, the maneuver utilizes the full width of the lane to minimize the curvature of the vehicle trace, thus reducing the lateral acceleration required to negotiate the turn. Third, the longitudinal velocity is reduced early in the maneuver such that the maximum tire forces are sufficient to negotiate the turn.

These coordinated adjustments to the motion plan are a direct result of the reduced tire force constraint set $\mathcal{P}_{k|t}^{\mathcal{U}}$, for all $k \in \{0, 1, \ldots, (N-1)\}$ in the optimization problem (12). The planned motion utilizes 90% of the available local traction, such that control authority is maintained throughout the maneuver. Thus, the planning/control scheme safely mitigates the critical situation. Next, we evaluate the traction adaptation concept in collision avoidance scenarios.

B. Scenario 2: Collision Avoidance at Low Local Traction

The second scenario is aimed at replicating the conditions when a pedestrian has suddenly entered the lane from the side of the road. At the start of the second scenario, $t_c = 0$ s, the pedestrian has been detected by the vehicle's perception stack and the planner/controller starts reacting. The initial velocity of the vehicle is 8 m/s and the pedestrian is detected 20 m ahead of the vehicle. Traction conditions are the same as in the previous case, that is, $\mu_{gt}(s) = 0.8$ for s < 0 and $\mu_{\rm gt}(s) = 0.2$ for $s \ge 0$. The NA scheme operates under static tire force constraints corresponding to a static friction estimate of $\mu_{sta} = 0.8$, while the adaptive scheme sets its tire force constraints dynamically according to $\mu_{est}(s) = \mu_{gt}(s)$. The objective of the vehicle is to avoid collision with the obstacle, while minimizing deviations from the lane center. The performance of the NA and TA schemes in this scenario is presented in Fig. 6.

1) Not Adapting to Low Local Traction: Under the nonadaptive scheme, we make the following observations. Upon detecting the pedestrian, Fig. 6(a), the NA scheme initiates an aggressive maneuver, that is, high curvature given the speed, that is meant to quickly move the vehicle away from its collision course with the obstacle. However, the vehicle is not able to realize the aggressive motion, so the plan is gradually shifted forward, ahead of the vehicle. At $t_c = 2$ s, Fig. 6(c), the vehicle has eventually moved 2 m to the right of the lane center, Fig. 6(j), enough to clear the obstacle. However, this occurs closer to the obstacle than was planned at $t_c = 0$ s. Also, at this point, the vehicle exhibits heavy under-steering behavior while turning right [see plotted front slip angle in Fig. 6(1)]. After clearing the obstacle, the vehicle starts steering left to return to the lane, but the grip is insufficient to counteract the motion out of the lane that was initiated during the first part of the maneuver. At $t_c = 3.1$ s, Fig. 6(e), the high slip angles have led to a complete loss of control authority, and the vehicle is now veering out of the lane. At this point in the experiment, the planner disengages and full braking is applied. The vehicle eventually comes to a stop 7 m outside the lane Fig. 6(g).

The underlying reason is again revealed by the planned tire forces at $t_c = 0$ s, plotted in Fig. 6(i). The NA scheme (solid blue) plans tire forces that exceed the physical limit (dotted blue), and therefore, the planned motion is not physically attainable. The fact that the real vehicle cannot realize the planned motion has two separate unfavorable consequences in this case: First, the vehicle cannot change direction as fast as the planner dictates, therefore, the margin to the obstacle ends up being smaller than what was originally planned [see Fig. 6(a) and (c)]. Second, the overly aggressive maneuver leads to high slip angles, Fig. 6(1), which eventually leads to complete loss of control authority, Fig. 6(e), due to tire saturation [4].

2) Adapting to Low Local Traction: As we saw in the previous scenario, the TA scheme dynamically updates tire force constraints such that the planned motions are dynamically feasible with respect to the local traction conditions. Therefore, in this case, we observe that the initial plan from the



Fig. 6. Scenario 2, collision avoidance at low traction conditions. (a)–(h) Overhead visualizations of the planned motion. The left and right columns correspond to the non-adaptive (NA) and traction adaptive (TA) schemes, respectively. (i) Planned tire forces and prevailing force limits for the front tire at $t_c = 0$ s. (j)–(m) Comparisons in terms of deviation from the lane center, forward velocity, front wheel slip angle, and lateral acceleration (a) NA, $t_c = 0$ s. (b) TA, $t_c = 0$ s. (c) NA, $t_c = 2$ s (under-steering). (d) TA, $t_c = 2$ s. (e) NA, $t_c = 3.1$ s (exiting lane). (f) TA, $t_c = 3.1$ s. (g) NA, $t_c = 5$ s (corresponding to (a) and (b)]. (j) Position w.r.t. lane center. (k) Forward velocity. (l) Front wheel slip angle. (m) Lateral acceleration.

TA scheme at $t_c = 0$ s, Fig. 6(b) is a less aggressive maneuver compared to the NA scheme. At $t_c = 2$ s, Fig. 6(b), the vehicle

has moved 2 m to the right of the lane center, enough to clear the obstacle, without generating excessive front wheel slip, Fig. 6(1). At $t_c = 3.1$ s, Fig. 6(f), the vehicle has cleared the obstacle and at $t_c = 5$ s, Fig. 6(h), the vehicle is returning from a comparatively small deviation from the lane, maintaining full control authority.

Fig. 6(i) reveals that the planned motion from the TA scheme at $t_c = 0$ s exploits approximately 90% ($\lambda = 0.90$) of the available traction force for large parts of the maneuver (dashed-dotted orange), while the NA scheme (solid blue) exceeds the traction bound by a large margin. Hence, apart from a small margin to the absolute traction limit, it delivers the most aggressive plan that is achievable at the prevailing conditions, yielding manageable slip angles, Fig. 6(1), and therefore maintained control authority. As such, the critical scenario is mitigated.

Furthermore, we observe from Fig. 6(j), showing deviation from the lane center, that for the first two seconds of the evasive maneuver, the sideways motion of the vehicle is almost identical for the two schemes. From this, we get an indication that sacrificing control authority of the vehicle in an overly aggressive maneuver does not improve the capacity to avoid colliding with the obstacle. It does, however, increase the risk of a secondary accident as the vehicle enters the opposing lane at high speed in an uncontrolled manner.

C. Scenario 3: Collision Avoidance at High Local Traction

A naive solution for handling scenarios 1 and 2 with static tire force constraints is to select a conservative static constraint, corresponding to the worst-case traction. However, this strategy will restrict the utilization of the vehicle's capability at favorable traction conditions. Next, we investigate this aspect for collision avoidance at high μ .

In this scenario, $\mu_{gt}(s) = 0.8$ for $s \ge 0$, and the pedestrian is detected 15 m ahead of the vehicle. The initial velocity of the vehicle is 15 m/s. The NA scheme operates under static tire force constraints corresponding to a conservative static friction estimate of $\mu_{sta} = 0.4$, while the adaptive scheme sets its tire force constraints dynamically according to $\mu_{est}(s) =$ $\mu_{gt}(s)$ as in previous scenarios. Performance for the NA and TA schemes in this scenario is presented in Fig. 7.

1) Not Adapting to High Local Traction: For the NA scheme, we make the following observations. Upon detecting the pedestrian, Fig. 7(a), the planner initiates an evasive maneuver under conservative static tire force constraints. Given the vehicle state at the time when the planner/controller starts reacting to the obstacle, in combination with the restrictive tire force constraints, there is no collision-free solution to the motion planning problem. As a result, the planner selects the least violating trajectory. At $t_c = 0.7$ s, Fig. 7(c), the vehicle is colliding with the obstacle at a velocity of 13.9 m/s. Notice that in this scenario the final motion of the vehicle corresponds well with the initial plan from the NA scheme, which was not the case in the low μ scenario. This is because in this case, the NA scheme is not violating the physical limitations in terms of tire forces [see Fig. 7(f) (blue dots)]. Rather, traction force utilization for the front wheel is consistently below 50% throughout the maneuver.



Fig. 7. Scenario 3, collision avoidance at high μ . (a)–(e) Overhead visualizations of the vehicle position and the planned motion. The left and right columns of overhead views correspond to the non-adaptive (NA) and traction adaptive (TA) schemes, respectively. (f) Planned tire forces and prevailing force limits for the front tire at $t_c = 0$ s. (g)–(j) Comparison in terms of deviation from the lane center, forward velocity, front wheel slip angle, and lateral acceleration. (a) NA, $t_c = 0$ s. (b) TA, $t_c = 0$ s. (c) NA, $t_c = 0.7$ s (accident). (d) TA, $t_c = 0.7$ s (evading). (e) TA, $t_c = 3$ s (no accident). (f) Planned tire forces at $t_c = 0$ s [corresponding to (a) and (b)]. (g) Position w.r.t. lane center. (h) Forward velocity. (i) Slip angle of front wheel. (j) Lateral acceleration.

2) Adapting to High Local Traction: Because $\mu_{est}(s) > \mu_{sta}$ in this case, adapting the tire force constraints means expanding the set of allowed tire forces compared to the NA case. Hence, the TA scheme is able to produce a more aggressive plan at $t_c = 0$ s, Fig. 7(b), that manages to avoid colliding with the obstacle. At $t_c = 0.7$ s, Fig. 7(d), the vehicle passes beside the obstacle as per the initial plan, and at $t_c = 3$ s, Fig. 7(e), the vehicle is returning to the lane, fully



Fig. 8. Minimum distances to avoid obstacles for three configurations of the planner: traction adaptive (TA) with $\mu_{est} = 0.8$ (blue), non-adaptive (NA) with static tire force constraints corresponding to $\mu_{sta} = 0.6$ (orange), and NA with static tire force constraints corresponding to $\mu_{sta} = 0.4$ (green). Zero value corresponds to collision.

recovered from the incident. Fig. 7(f) reveals that at $t_c = 0$ s, the TA scheme plans to exploit a larger part of the available traction. This is also reflected by the measured front slip angle, Fig. 7(i).

Also, in this scenario, because of the higher accelerations, the effect of dynamic normal loads become more apparent, see Fig. 7(f). In the beginning of the maneuver, when the vehicle is braking hard, the normal force, and hence the maximum horizontal force, is larger on the front tire and smaller on the rear. Notice that the TA scheme is able to exploit also this effect and allocates 12% higher front tire force to the front tire and the start of the maneuver (while the vehicle is braking hard) compared to the end of the maneuver (when the vehicle is no longer decelerating). Although this effect is small compared to the 50% difference between the TA and NA schemes due to the friction estimate, it is not insignificant. The effect will be more pronounced at higher accelerations, and thus, is even more important to consider for vehicles that are smaller and more agile than our 8.5 ton test vehicle.

The higher utilization of the available traction for the TA scheme leads to a quicker sideways motion compared to the NA scheme [Fig. 7(g)], and for this case, such a performance difference represents the difference between success and failure of the collision avoidance maneuver. Here, in contrast to Scenarios 1 and 2, the TA scheme mitigates the critical situation by expanding the tire force constraints instead of reducing them. In all three cases, however, it utilizes around 90% of the locally available traction.

From this example, we conclude that conservative friction estimation may reduce collision avoidance performance. To further investigate the extent of this phenomenon, we ran three configurations of the planning/ctrl schemes, TA with $\mu_{est} = 0.8$, NA with $\mu_{sta} = 0.6$, and NA with $\mu_{sta} = 0.4$, through a sequence of 40 challenging collision avoidance scenarios each. A sequence of relative positions of the suddenly appearing obstacles was pre-generated by sampling from a uniform distribution, $-1.0 \text{ m} \le d_{obs} \le 1.0 \text{ m}$, $s + 11.0 \text{ m} \le s_{obs} \le s + 13.0 \text{ m}$. The initial velocity at obstacle detection was 10 m/s.

With the three configurations, the vehicle managed to avoid collision in 11/40, 6/40, and 4/40 cases, respectively. Fig. 8 shows individual differences in terms of minimum distance

margin to the obstacle, Δ_m , for the 11 cases in which at least one configuration performed a successful avoidance maneuver.

Note that the most conservative configuration (NA with $\mu_{sta} = 0.4$) has the least amount of successful avoidance maneuvers and the four successful avoidance maneuvers are very near misses, with less than 0.10 m of clearance to the obstacle. The least conservative configuration (TA with $\mu_{est} = 0.8$) has the largest number of successful avoidance maneuvers and tends to also have the largest amount of clearance to the avoided obstacles. This is true for all cases except for obstacles 5 and 8, where (NA with $\mu_{sta} = 0.6$) has slightly higher margin. We believe this stems from variations in the state of the vehicle when the obstacle is detected. In particular, the lateral position within the lane exhibited small variations during the experiment, which had noticeable effect on the margin to the obstacle.

D. Scenario 4: Collision Avoidance With Discrete Decisions

As mentioned in Section IV-D, the RTI method handles situations when the initial guess \hat{T}_t violates state constraints, $\hat{x}_{k|t} \notin \mathcal{P}_{k|t}^{\mathcal{X}}$, for any $k \in \{0, 1, \dots, N\}$, by softening the constraints and giving constraint violations a high cost. This in combination with the fact that the algorithm only searches for a solution locally around \hat{T}_t leads to sensitivity to highly suboptimal local minima. The problem occurs when the constraint configuration changes abruptly between planning iterations, for example, at suddenly appearing obstacles, or in our approach to traction adaptation, at large changes in traction conditions. Therefore, we employ the sampling augmentation procedure, see Sections V-B and V-C, to mitigate this problem. In this section, we evaluate the safety implications of sampling augmentation in critical scenarios with discrete decisions.

Scenario 4 is a collision avoidance scenario where two obstacles appear simultaneously 20 and 35 m in front of the vehicle. The first obstacle appears slightly to the left of the lane center and the second obstacle appears on the right side of the lane. State constraints are set such that obstacles should be avoided by a margin of 0.5 m. We compare performance of the non-augmented RTI scheme (abbreviated RTI), with our sampling augmented variant (abbreviated SA). For this test, both schemes use the same static tire force constraints. Results are presented in Fig. 9.

1) Without Sampling Augmentation (*RTI*): When the obstacles are detected at $t_c = 1.0$ s, Fig 9(a), the non-augmented RTI scheme produces an evasive plan that avoids the obstacles by going in between them (right of the first and left of the second obstacle). The vehicle manages to avoid collision, Fig. 9(c) and (e), but needs to use the maximum front tire force almost throughout the maneuver, Fig. 9(g), and has a very small distance margin Δ_m to the obstacles (state constraints are violated), Fig. 9(h).

The decision to plan the maneuver between the obstacles stems from the last planning iteration before the obstacles appeared. At the first planning iteration after the obstacle is detected, the solution from the previous iteration $\mathcal{T}_{t-1}^{\star}$ and hence the initial guess, $\hat{\mathcal{T}}_t$, is at the center of the lane. When the obstacles appear, $\hat{\mathcal{T}}_t$ violates the state constraints $\hat{x}_{k|t} \notin \mathcal{P}_{k|t}^{\mathcal{X}}$



Fig. 9. Scenario 4, collision with discrete decisions. (a)–(f) Overhead visualizations of the vehicle position and the planned motion. The left and right columns of overhead views correspond to the non-augmented RTI scheme (RTI) and sampling augmented (SA) schemes, respectively. (g) Planned tire forces for the front tires at $t_c = 1.0$ s. (h) Comparison in terms of smallest distance to either obstacle. (i) Which candidate was selected as initial guess with the SA scheme. (a) RTI, $t_c = 1$ s. (b) SA, $t_c = 4.5$ s. (c) RTI, $t_c = 2.5$ s. (d) SA, $t_c = 2.5$ s. (e) RTI, $t_c = 5.2$ s. (f) SA, $t_c = 4.5$ s. (g) Planned tire forces at $t_c = 1.0$ s (c) corresponding to (a) and (b)]. (h) Distance to obstacles. (i) Initial guess selection with SA.

for some $k \in \{0, 1, ..., N\}$. Since state constraints are soft, the optimization problem is still feasible, but the decision of adjusting the trajectory left or right of an obstacle is determined by minimizing cost locally around $\mathcal{T}_{t-1}^{\star}$. Since in this case, the first obstacle appears slightly to the left of the lane center, the planner makes the decision to go between the obstacles.

2) With Sampling Augmentation (SA): When the sampling augmented RTI algorithm is exposed to the same scenario, we observe that upon detecting the obstacle, the SA scheme instead produces an evasive plan to the left of both obstacles, Fig. 9(b). When following this plan, the vehicle is able to clear the obstacles more comfortably, Fig. 9(c) and (e). It uses the maximum front tire force for a smaller amount of the time, Fig. 9(g), keeps above the minimum margin to both obstacles (not violating state constraints), Fig. 9(h), and recovers more quickly after the incident [compare Fig. 9(e) and (f)].

Just as before, at the first planning iteration after the obstacles are detected, the initial guess \tilde{T}_t originating from $\mathcal{T}_{t-1}^{\star}$ is in collision with the obstacles. However, due to the sampling augmentation procedure, at $t_c = 0.1$ s, a sampled trajectory \hat{T}_t' is selected as initial guess. In the subsequent planning iterations throughout the maneuver, \tilde{T}_t is selected, Fig. 9(i).

Fig. 9(i) reveals that the sampled trajectories were only used in a single planning iteration just after the obstacles appeared. This indicates that sampling augmentation steps in exactly when needed, to avoid unfavorable local minima, that is, at rapid changes in the constraint configuration of the motion planning problem. Otherwise, when changes to the constraint configuration are small, \tilde{T}_t typically has the lowest cost and gets selected. This result indicates that sampling augmentation alleviates Problem 3 of Section IV-D RTI, without impairing the otherwise desirable properties of the RTI method.

The sampling augmentation procedure should be viewed as an exploratory effort, which continuously evaluates a range of dynamically feasible maneuver alternatives, spread throughout the physically reachable drivable area, such that discrete decision-making is made based on the full reachable state space of the vehicle (given current traction conditions), represented by $\hat{S}_t \cup \tilde{T}_t$ rather than being solely determined by local information around T_{t-1}^* .

VII. CONCLUSION AND FUTURE WORK

In this article, we address the problem of motion planning at the limits of handling under locally varying traction conditions. We evaluate the proposed traction adaptive algorithm by comparing it to an equivalent non-adaptive scheme in a sequence of critical scenarios at various traction conditions. Results from experimental evaluation indicate that traction adaptation improves capacity to avoid accident in all four tested critical scenarios. The proposed algorithm's performance in this regard stems from two properties.

- 1) Ensured dynamic feasibility of planned motions, even at varying local traction conditions.
- A high ratio (90% for the tested configuration) of locally available traction is utilized to avoid collision if necessary.

Moreover, we demonstrate experimentally that the sampling augmentation procedures (steps A–D of Section V) of the proposed algorithm successfully mitigates the computational challenges associated with infeasibility and sensitivity to local minima (listed in Section IV-D) that emerges in the planning problem at rapid changes to the constraint configuration.

In Section III-D, we briefly introduce a selection of approaches to obtaining predictive friction estimates online. It is intuitively clear that the characteristics and performance of such friction estimates will influence the planned behavior. We do not include that aspect in the scope of this article, instead we present a follow-up study, dedicated to that topic [39], where we conclude that combining camera-based and traditional friction estimation techniques enable near-optimal traction adaptive motion planning and control. These preliminary results are obtained by emulating

TABLE I STATIC PHYSICAL PARAMETERS OF TEST VEHICLE

Parameter	Value
m	8350 kg
I_z	8150 kgm^2
h	1.0 m
$l_{\rm f}$	1.20 m
$l_{\rm r}$	2.20 m

state-of-the-art friction estimation methods in a simulated environment. Thus, the next step is to return to the test track and evaluate traction adaptive motion planning and control in closed loop with state-of-the-art friction estimation.

APPENDIX A EXPERIMENTAL SETUP

The vehicle used for the experiments, depicted in Fig. 1, was a Volvo FH16 6 \times 4, a 750 hp tractor equipped with Volvo Dynamic Steering and I-Shift transmission, housed at the vehicle research laboratory Resource for Vehicle Research (REVERE) at the Chalmers University of Technology. Table I outlines the physical parameters of the vehicle. The truck is equipped with multiple sensors to perceive the vehicle's state and surroundings, including LiDAR, stereo vision, differential GNSS, and IMU.

The computing system of the truck includes the microservices-based software stack OpenDLV⁴ that enables transparent communication of distributed software components using inter-process communication (IPC), for example, network communication and shared memory for low latency in real-time setups and for large volume sensor data handling.

The planning/control algorithms were integrated in a ROS stack of rudimentary autonomous functionality (localization, state estimation, track and object handling, etc.) which was run on a PC with an Intel Core i7-7820HK CPU @ 2.9 GHz and an Nvidia GeForce GTX 1070 GPU. On the truck's computing node, a ROS-to-OpenDLV bridge was realized for the project to enable the ROS stack to access of the truck's sensors and actuators. Experiments were primarily conducted at the High Speed Area of the AstaZero test facility in Borås, Sweden, with complementary low μ tests at the Stora Holm test track in Gothenburg, Sweden.

APPENDIX B

DYNAMIC BICYCLE MODEL WITH TIME-VARYING TIRE FORCE LIMITS IN ROAD ALIGNED COORDINATES

We derive a dynamic bicycle model based on standard vehicle dynamics literature [4]. See Fig. 10 for the geometric relations of the model. From Newton's second law for the lateral dimension, we have that

$$\dot{v}_y = \frac{1}{m}(F_{yf} + F_{yr}) - v_x \dot{\psi} + g \sin \phi$$

where F_{yf} and F_{yr} are lateral forces on the front and rear tires, respectively. The term $v_x \dot{\psi}$ is the cross product term

⁴Cf. https://github.com/chalmers-revere/opendlv

from the rotating vehicle-fixed coordinate system⁵ and $g \sin \phi$ is the contribution from the bank angle of the road. As for the longitudinal dynamics, we have that

$$\dot{v}_x = \frac{1}{m}(F_{xf} + F_{xr}) - g\sin(\theta)$$

where F_{xf} and F_{xf} are the controlled tractive forces on the front and rear tires, respectively. The term $g\sin(\theta)$ represent the longitudinal acceleration contribution due to the inclination of the road. The yaw dynamics are given by moment balance around the *z*-axis as

$$\ddot{\psi} = \frac{1}{I_z} \left(l_{\rm f} F_{\rm yf} - l_{\rm r} F_{\rm yr} \right)$$

where I_z is the vehicle's moment of inertia about the *z*-axis, and l_f and l_f are the distances from the center of mass to the front and rear axle, respectively. In order to capture the motion of the vehicle relative to the road, we use a road aligned coordinate frame [40]. A coordinate *s* represents the progress along the centerline of the lane and a coordinate *d* represents the lateral deviation from the centerline, as depicted in Fig. 10(a). The *d*-coordinate increases to the left. Variables ψ_c and κ_c denote the path tangent and path curvature at the point on the centerline, respectively, where the vehicle position is perpendicular to the centerline tangent. $\Delta \psi = \psi - \psi_c$ represents the relative angle between the vehicle yaw angle and the tangent of the centerline. Geometric relations of the coordinate system definition give us the evolution of *s*, *d*, and $\Delta \psi$. Finally, we express the continuous time dynamics as

$$\dot{s} = \frac{v_x \cos(\Delta \psi) - v_y \sin(\Delta \psi)}{1 - d\kappa_c}$$
$$\dot{d} = v_x \sin(\Delta \psi) + v_y \cos(\Delta \psi)$$
$$\Delta \dot{\psi} = \dot{\psi} - \kappa_c \frac{v_x \cos(\Delta \psi) - v_y \sin(\Delta \psi)}{1 - d\kappa_c}$$
$$\ddot{\psi} = \frac{1}{I_z} (l_f F_{yf} - l_r F_{yr})$$
$$\dot{v}_x = \frac{1}{m} (F_{xf} + F_{xr}) - g \sin(\theta)$$
$$\dot{v}_y = \frac{1}{m} (F_{yf} + F_{yr}) - v_x \dot{\psi} + g \sin \phi.$$
(14)

By selecting the state and control vectors as $x = [s, d, \Delta \psi, \dot{\psi}, v_x, v_y]^{\top}$ and $u = [F_{yf}, F_{xf}, F_{xr}]^{\top}$, we can compactly write the planning model (14) as $\dot{x} = f_p(x, u)$.

The rear lateral force F_{yr} is not included among the control inputs, instead it is computed from the state through a time-varying linear tire model

$$F_{\rm yr} = \tilde{C}_{\rm r}(\mu) \alpha_{\rm r} \tag{15}$$

with the rear slip angle $\alpha_r = -\arctan((v_y - l_r \dot{\psi}))/v_x$. The time-varying rear cornering stiffness \tilde{C}_r is obtained by linearizing Pacejka's Magic Formula tire model [41]

$$F_{y} = D(\mu) \sin(C(\mu) \arctan(B(\mu)\alpha_{f} + \dots - E(\mu)(B(\mu)\alpha_{f} - \arctan(B(\mu)\alpha_{f}))))$$

⁵Since v_y is typically small, we ignore the corresponding term $v_y \psi$ in the equation for v_x .







Fig. 10. Geometric relations of the dynamic bicycle model with time-varying tire force limits in a road aligned coordinate frame. (a) Top down view. (b) Left side view.

with parameters $B(\mu)$, $C(\mu)$, $D(\mu)$, and $E(\mu)$ obtained as functions of μ from a lookup table that approximately associate the tire parameters to the friction coefficient for various road surface conditions.

Remark 5: In general, the linear tire model used in (15) does not accurately predict the relation between slip angle and tire force at high slip angles [4]. However, in this context, the combined lateral and longitudinal force command is always inside the friction circle, by a margin determined by the utilization factor λ . Inside such a boundary, a time-varying linear model corresponds well with the Magic Formula. Considering the computational trade-off mentioned in Section II, this design represents a compromise between complexity and accuracy that have proved to be practically viable for our application.

The normal forces F_{zf} and F_{zr} acting on the tires of the model are obtained from moment balances about the contact points of the rear and front tires respectively as

$$F_{zf} = \frac{1}{l_f + l_r} \left(-m\dot{v}_x h - \text{mgh}\sin\theta + \text{mgl}_r\cos\theta \right)$$

$$F_{zr} = \frac{1}{l_f + l_r} \left(m\dot{v}_x h + \text{mgh}\sin\theta + \text{mgl}_f\cos\theta \right).$$
(16)

See geometric relations in Fig. 10(b). Combining these relations with the friction estimate μ , we define front and rear *time-varying*⁶ force boundaries as

$$F_{\rm f}^{\rm (ub)} = \lambda \mu F_{z\rm f}, \quad F_{\rm r}^{\rm (ub)} = \lambda \mu F_{z\rm r}$$

⁶Variables μ , F_{zf} , and F_{zr} vary in time and hence, so do $F_{f}^{(ub)}$ and $F_{r}^{(ub)}$. Just as with the state variables in (14), the time dependency is omitted from the equations to simplify notation.



Fig. 11. Input constraint polytope $\mathcal{P}^{\mathcal{U}}$ (for single prediction step *k*) bounding control signals produced by the algorithm (dark gray) and theoretical traction limits according to (1).

where $\lambda \in [0, 1]$ is a user adjustable traction utilization factor. If $\lambda = 1.0$, the algorithm is allowed to utilize the full friction circle. The experiments for this article were run at $\lambda = 0.9$. To comply with these bounds, tire force commands need to fulfill

$$\sqrt{F_{xf}^2 + F_{yf}^2} \le F_f^{(ub)}, \quad \sqrt{F_{xr}^2 + F_{yr}^2} \le F_r^{(ub)}$$
 (17)

where F_{yr} is determined by (15). In addition to the bounded tire forces (17), the drivetrain of the vehicle imposes additional constraints

$$F_{xf} \leq 0, \quad F_{xr} \leq F_r^{(ub,dr)}$$
 (18)

due to rear wheel drive and limited tractive force from the motor $F_r^{(ub,dr)}$.

Finally, the algorithmic framework used in this article requires linear input constraints. To accommodate this, we under-approximate the input constraints (17) and (18) by a 3-D convex polytope $\mathcal{P}^{\mathcal{U}}$, specific to each prediction step k, visualized in Fig. 11.

APPENDIX C Low-Level Control Interface

The purpose of the control interface is to translate tire force commands, F_{yf} , F_{xf} , and F_{xr} from the planned trajectory \mathcal{T}_t^{\star} into the actual control inputs of the vehicle. For our particular test vehicle, those are steering angle request δ_{req} and longitudinal acceleration request a_{req} .

For the longitudinal control input, the transformation is trivially computed as $a_{req} = (F_{xf} + F_{xr})/m$. For the steering input δ_{req} , we begin by computing the kinematic component of the steering angle, given by δ_{kin} , that is, the steering angle that would yield the planned motion in the absence of tire slip. We fit a circle segment of curvature ρ to the Cartesian representation of the position states s_{klr}^* and d_{klr}^* for $k \in$ $\{0, 1, \ldots, (N_{fit})\}$ of \mathcal{T}_t^* , where N_{fit} is selected such that only the initial part of the trajectory is used. Then, the kinematic steering angle component is obtained as $\delta_{kin} = \rho(l_f + l_r)$. The dynamic steering angle component, that is, the desired front wheel slip angle, is obtained as $\alpha_f = F_{yf}/\tilde{C}_f$, where the time-varying cornering stiffness \tilde{C}_f is obtained by linearizing the Pacejka Magic Formula tire model as per the procedure described in Appendix B. The steering angle request is then obtained as $\delta_{\text{req}} = \delta_{\text{kin}} + \alpha_{\text{f}}$. Finally, the control commands δ_{req} and a_{req} are passed to the vehicle via the ROS-to-OpenDLV bridge described in Appendix A.

ACKNOWLEDGMENT

Disclaimer: The document reflects only the author's view and the Commission is not responsible for any use that may be made of the information it contains.

REFERENCES

- A. Liniger, A. Domahidi, and M. Morari, "Optimization-based autonomous racing of 1: 43 scale RC cars," *Optim. Control Appl. Methods*, vol. 36, no. 5, pp. 628–647, 2015.
- [2] J. K. Subosits and J. C. Gerdes, "From the racetrack to the road: Realtime trajectory replanning for autonomous driving," *IEEE Trans. Intell. Vehicles*, vol. 4, no. 2, pp. 309–320, Jun. 2019.
- [3] D. Fambro, K. Fitzpatrick, and R. Koppa, NCHRP Report 400: Determination of Stopping Sight Distances. Washington, DC, USA: TRB, National Research Council, 1997.
- [4] R. Rajamani, Vehicle Dynamics and Control. New York, NY, USA: Springer, 2011.
- [5] J. P. Alsterda, M. Brown, and J. C. Gerdes, "Contingency model predictive control for automated vehicles," in *Proc. Amer. Control Conf.* (ACC), Jul. 2019, pp. 717–722.
- [6] L. Svensson, M. Bujarbaruah, N. R. Kapania, and M. Torngren, "Adaptive trajectory planning and optimization at limits of handling," in *Proc. IEEE/RSJ Int. Conf. Intell. Robots Syst. (IROS)*, Nov. 2019, pp. 3942–3948.
- [7] B. Paden, M. Čáp, S. Z. Yong, D. Yershov, and E. Frazzoli, "A survey of motion planning and control techniques for self-driving urban vehicles," *IEEE Trans. Intell. Veh.*, vol. 1, no. 1, pp. 33–55, Mar. 2016.
- [8] D. González, J. Pérez, V. Milanés, and F. Nashashibi, "A review of motion planning techniques for automated vehicles," *IEEE Trans. Intell. Transp. Syst.*, vol. 17, no. 4, pp. 1135–1145, Apr. 2016.
- [9] J. Ziegler, P. Bender, T. Dang, and C. Stiller, "Trajectory planning for Bertha—A local, continuous method," in *Proc. IEEE Intell. Vehicles Symp.*, Jun. 2014, pp. 450–457.
- [10] S. Boyd and L. Vandenberghe, *Convex Optimization*. Cambridge, U.K.: Cambridge Univ. Press, 2004.
- [11] J. T. Betts, "Survey of numerical methods for trajectory optimization," J. Guid. Control Dyn., vol. 21, no. 2, pp. 193–207, Mar./Apr. 1998.
- [12] C. R. Hargraves and S. W. Paris, "Direct trajectory optimization using nonlinear programming and collocation," J. Guid., Control, Dyn., vol. 10, no. 4, pp. 338–342, 1987.
- [13] O. von Stryk and R. Bulirsch, "Direct and indirect methods for trajectory optimization," Ann. Oper. Res., vol. 37, no. 1, pp. 357–373, 1992.
- [14] M. Diehl, H. G. Bock, and J. P. Schlöder, "A real-time iteration scheme for nonlinear optimization in optimal feedback control," *SIAM J. Control Optim.*, vol. 43, no. 5, pp. 1714–1736, Jul. 2005.
- [15] M. Werling and D. Liccardo, "Automatic collision avoidance using model-predictive online optimization," in *Proc. IEEE 51st Annu. Conf. Decision Control (CDC)*, Dec. 2012, pp. 6309–6314.
- [16] T. Howard and C. Green, "State space sampling of feasible motions for high-performance mobile robot navigation in complex environments," *J. Field Robot.*, vol. 25, no. 1, pp. 325–345, Jun. 2008.
- [17] D. Ferguson, T. M. Howard, and M. Likhachev, "Motion planning in urban environments," *J. Field Robot.*, vol. 25, nos. 11–12, pp. 939–960, Nov. 2008.
- [18] M. Werling, S. Kammel, J. Ziegler, and L. Gröll, "Optimal trajectories for time-critical street scenarios using discretized terminal manifolds," *Int. J. Robot. Res.*, vol. 31, no. 3, pp. 346–359, 2012.
- [19] M. Tanaskovic, L. Fagiano, R. Smith, and M. Morari, "Adaptive receding horizon control for constrained MIMO systems," *Automatica*, vol. 50, no. 12, pp. 3019–3029, Dec. 2014.
- [20] L. Hewing, J. Kabzan, and M. N. Zeilinger, "Cautious model predictive control using Gaussian process regression," *IEEE Trans. Control Syst. Technol.*, vol. 28, no. 6, pp. 2736–2743, Nov. 2020.
- [21] J. Kohler, E. Andina, R. Soloperto, M. A. Müller, and F. Allgower, "Linear robust adaptive model predictive control: Computational complexity and conservatism," in *Proc. IEEE 58th Conf. Decis. Control (CDC)*, Dec. 2019, pp. 1383–1388.
- [22] F. Gustafsson, "Slip-based tire-road friction estimation," Automatica, vol. 33, no. 6, pp. 1087–1099, Jun. 1997.

- [23] R. Rajamani, G. Phanomchoeng, D. Piyabongkarn, and J. Y. Lew, "Algorithms for real-time estimation of individual wheel tire-road friction coefficients," *IEEE/ASME Trans. Mechatronics*, vol. 17, no. 6, pp. 1183–1195, Dec. 2012.
- [24] J. Yi, L. Alvarez, and R. Horowitz, "Adaptive emergency braking control with underestimation of friction coefficient," *IEEE Trans. Control Syst. Technol.*, vol. 10, no. 3, pp. 381–392, May 2002.
- [25] L. Alvarez, J. Yi, R. Horowitz, and L. Olmos, "Adaptive emergency braking control with observer-based dynamic tire/road friction model and underestimation of friction coefficient," *IFAC Proc. Volumes*, vol. 35, no. 1, pp. 211–216, 2002.
- [26] P. Falcone, F. Borrelli, J. Asgari, H. E. Tseng, and D. Hrovat, "Predictive active steering control for autonomous vehicle systems," *IEEE Trans. Control Syst. Technol.*, vol. 15, no. 3, pp. 566–580, May 2007.
- [27] Y. Gao, T. Lin, F. Borrelli, E. Tseng, and D. Hrovat, "Predictive control of autonomous ground vehicles with obstacle avoidance on slippery roads," in *Proc. ASME Dyn. Syst. Control Conf.*, vol. 1, Jan. 2010, pp. 265–272.
- [28] S. Roychowdhury, M. Zhao, A. Wallin, N. Ohlsson, and M. Jonasson, "Machine learning models for road surface and friction estimation using front-camera images," in *Proc. Int. Joint Conf. Neural Netw. (IJCNN)*, Jul. 2018, pp. 1–8.
- [29] M. Nolte, N. Kister, and M. Maurer, "Assessment of deep convolutional neural networks for road surface classification," in *Proc. 21st Int. Conf. Intell. Transp. Syst. (ITSC)*, Nov. 2018, pp. 381–386.
- [30] F. Holzmann, M. Bellino, R. Siegwart, and H. Bubb, "Predictive estimation of the road-tire friction coefficient," in *Proc. IEEE Int. Conf. Control Appl.*, Oct. 2006, pp. 885–890.
- [31] R. Omer and L. Fu, "An automatic image recognition system for winter road surface condition classification," in *Proc. 13th Int. IEEE Conf. Intell. Transp. Syst.*, Sep. 2010, pp. 1375–1379.
- [32] Y. Qian, E. J. Almazan, and J. H. Elder, "Evaluating features and classifiers for road weather condition analysis," in *Proc. IEEE Int. Conf. Image Process. (ICIP)*, Sep. 2016, pp. 4403–4407.
- [33] L. Chisci, J. A. Rossiter, and G. Zappa, "Systems with persistent disturbances: Predictive control with restricted constraints," *Automatica*, vol. 37, no. 7, pp. 1019–1028, 2001.
- [34] W. Langson, I. Chryssochoos, S. V. Raković, and D. Q. Mayne, "Robust model predictive control using tubes," *Automatica*, vol. 40, no. 1, pp. 125–133, 2004.
- [35] M. Bujarbaruah, U. Rosolia, Y. R. Sturz, and F. Borrelli, "A simple robust MPC for linear systems with parametric and additive uncertainty," in *Proc. Amer. Control Conf. (ACC)*, May 2021, pp. 2108–2113.
- [36] S. Gros, M. Zanon, R. Quirynen, A. Bemporad, and M. Diehl, "From linear to nonlinear MPC: Bridging the gap via the real-time iteration," *Int. J. Control*, vol. 93, no. 1, pp. 62–80, 2016.
- [37] B. Houska, H. J. Ferreau, and M. Diehl, "An auto-generated realtime iteration algorithm for nonlinear MPC in the microsecond range," *Automatica*, vol. 47, no. 10, pp. 2279–2285, 2011.
- [38] E. C. Kerrigan and J. M. Maciejowski, "Soft constraints and exact penalty functions in model predictive control," in *Proc. UKACC Int. Conf. Control*, 2000, pp. 1–6.
- [39] L. Svensson and M. Törngren, "Fusion of heterogeneous friction estimates for traction adaptive motion planning and control," 2021, arXiv:2105.06692.
- [40] A. Micaelli and C. Samson, "Trajectory tracking for unicycle-type and two-steering-wheels mobile robots," INRIA, Tech. Rep. 2097, 1993.
- [41] H. B. Pacejka and E. Bakker, "The magic formula tyre model," Vehicle Syst. Dyn., vol. 21, no. S1, pp. 1–18, 1992.



Lars Svensson (Member, IEEE) received the M.Tech. degree in engineering physics from Uppsala University, Uppsala, Sweden, in 2015, and the Ph.D. degree in mechatronics from the KTH Royal Institute of Technology, Stockholm, Sweden, in 2021.

In Fall 2018, he was a Visiting Researcher with the Model Predictive Control Lab, University of California at Berkeley, Berkeley, CA, USA. His research focuses on motion planning and control of mobile robots in challenging operational conditions, including road vehicles at the limits of handling and heavy duty vehicles in rough terrain.



Monimoy Bujarbaruah received the M.Tech. degree in aerospace engineering from the Indian Institute of Technology (IIT) Bombay, Mumbai, India, in 2016. He is currently pursuing the Ph.D. degree in mechanical engineering with the University of California at Berkeley, Berkeley, CA, USA. He is also a Graduate Student Researcher with

Model Predictive Control (MPC) Lab, under Prof. Francesco Borrelli. His research interests include adaptive and optimal control, learning-based algorithms for predictive control, convex optimization,

and embedded control with primary applications to autonomous vehicles and driver assistance systems.



Christian Berger received the Ph.D. degree from RWTH Aachen University, Aachen, Germany, in 2010.

He is currently an Associate Professor with the Department of Computer Science and Engineering, University of Gothenburg, Gothenburg, Sweden. He coordinated the research project for the vehicle "Caroline," which successfully participated in the 2007 DARPA Urban Challenge. He also co-led the Chalmers Truck Team during the 2016 Grand Cooperative Driving Challenge (GCDC) and is one

of the two leading architects behind OpenDLV. His research expertise is on distributed real-time software, microservices for embedded systems and cyberphysical systems, and continuous integration/deployment/experimentation for embedded systems.



Arpit Karsolia received the master's degree in automotive systems from the HAN University of Applied Sciences, Arnhem, The Netherlands, in 2014.

Since 2015, he has been working with the Chalmers University of Technology, Gothenburg, Sweden, in a research lab called REVERE. From 2015 to 2017, he was hired as a Consultant at the lab in the role of Lab Technician. Since 2017, he has been working as a Research Engineer at the lab, focusing on building and testing of autonomous light and heavy vehicles.



Martin Törngren (Senior Member, IEEE) received the Ph.D. degree in mechatronics from the KTH Royal Institute of Technology, Stockholm, Sweden, in 1995.

He has been a Professor in embedded control systems with the Division of Mechatronics and Embedded Control Systems, KTH Royal Institute of Technology, since 2002. He has been a pioneer in bridging the gaps between the fields of automatic control and real-time distributed systems, with research in recent years focused on architectural

design and safety of autonomous trustworthy cyber-physical systems. He has spent several periods abroad as a visiting scholar, including at UC Berkeley, Berkeley, CA, USA, and the Stevens Institute of Technology, Hoboken, NJ, USA. Networking and multidisciplinary research have been characteristic throughout his career. He created the Innovative Center for Embedded Systems, a KTH-industry competence network launched in 2008, and he is the Director of the Swedish research center on Trustworthy Edge Computing Systems and Applications, TECoSA.